Discovering and Visualizing Hierarchy in the Data

Kun Yang, Wing Hung Wong

Abstract—How to extract useful insights from data in a human perceivable manner is always a challenge when the dimension and amount of the data is large. Often, the data can be organized according to certain hierarchical structure that are stemmed either from data collection process or from the information and phenomena carried by the data itself. The current study attempts to discover and visualize these underlying hierarchies. Regarding each observation as a draw from a (hypothetical) multidimensional joint density, our first goal is to approximate this unknown density with a piecewise constant function over the binary partitioned sample space; our non-parametric approach makes no assumptions on the form of the density, such as assuming that it is Multivariate Gaussian, or that it is a mixture of a small number of Gaussians. Given the piecewise constant density function and its corresponding partitions of the sample space, our second goal is to construct a connected graph and build up a tree representation of the data from sub-level sets. To demonstrate that our method is a general data mining and visualization tool which can provide "multi-resolution" summaries and reveal different levels of information of the data, we apply it to two real data sets from different fields.

Index Terms—Visualization, Hierarchy, Binary Partition, Discrepancy

arXiv:1403.4370v3 [stat.AP] 28 Mar 2014

1 Introduction

With the advent of "Big Data" era, massive amount of data are collected and processed at an unprecedented speed. The explosion of data has been exploited in almost every field. It supports online business in commercial world, disease diagnosis and new drug development in pharmaceutical industry, governmental effort in countering terrorism and cutting-edge research in academia. The data are typically featured with high dimension and large sample size. Instead of interpreting the data through specific models, how to understand them in a human perceivable manner is a challenge. Besides the traditional tables, maps and diagrams, advanced data visualization techniques are used to reveal dynamic information in systems [18, 19], discover connections in networks [7] and extract topology in data clouds [10].

Many data manifest certain patterns of hierarchies that are stemmed either from data collection process or from the information and phenomena carried by the data itself. Examples include census data collected at county level, state level or nation level and stem cell population differentiated into various specialized cell types. In this paper, we propose new algorithms to discover and visualize these underlying hierarchies. We first introduce the concept of binary partitions; then develop the method to construct a class of piecewise constant density estimator, by regarding the multidimensional data as independent observations generated from some hypothetical distribution; the method is motivated by the discrepancy criteria in Quasi Monte Carlo and has worst complexity $O(n\log^d n)$, where d is the dimension and n is the sample size. Subsequently, we introduce the tree of sub-level sets and the algorithm to build it according to the density estimate. Through simulation and real data examples, it is shown that this binary partition based density estimate and its corresponding tree of sub-level sets provide a general tool to mine and visualize data and are capable of revealing and summarizing hidden hierarchical structures.

2 Binary Partition by Discrepancy

Let Ω be a hypercube in $\mathbb{R}^d$. A binary partition P on Ω is a collection of sub-cubes whose union is Ω. Starting with $P_1 = \{\Omega\}$ at level 1 and $P_t = \{\Omega_1, \Omega_2, ..., \Omega_t\}$ at level t, P_{t+1} is produced by dividing one of regions in P_t into two sub-cubes along one of its coordinates, then combining these two sub-cubes with the rest of regions in P_t; continuing with this fashion, one can generate any binary partition in any level (Figure 1).

Piecewise constant function is of fundamental importance in mathematics and statistics for its simplicity and its ability to approximate any continuous function to any degree. In order to construct a simple yet flexible density estimator, we restrict the class of density function as the piecewise constant function on the binary partitioned sample space. Our algorithm, by exploiting the sequential build-up of binary partition, can find an optimal density estimation efficiently.

To decide whether a sub-cube deserves further partitioning, we need to test the uniformity of points in it. Since any sub-cube is a translation and scaling of unit cube, it is equivalent to test the following hypothesis,

$$H_0 : x \sim U[0,1)^d, x \in \mathscr{S} = \{x_i = (x_{i1}, ..., x_{id}), x_i \in [0,1)^d\}_{i=1}^n$$

In the literature of quasi-Random Number Generators or quasi-Monte Carlo methods [9], there are a number of criteria for measuring whether a set of points is uniformly scattered in the unit cube $[0,1)^d$. These criteria are called discrepancies, and they arise in the error analysis of quasi-Monte Carlo methods for evaluating multiple integrals [14].

The precise definitions of the discrepancy and the variation depend on the space of integrands. The one widely used in quasi-Monte Carlo analysis is the classic star discrepancy, defined as follows,

$$D^*(\mathscr{S}) = \sup_{x \in [0,1)^d} \left| \frac{\#(\mathscr{S} \cap [0,x))}{n} - \prod_{j=1}^d x_j \right|$$

where # is the cardinality of a set. Besides D^*, there are $\mathscr{L}^2$ star discrepancy, symmetric discrepancy and centered discrepancy defined on the reproducing kernel Hilbert space, they all have interesting geometrical interpretations and one of the advantages is that their explicit formulas are available [8]. Thus, we can construct useful statistics for testing multivariate uniformity on a set of points according to these formulas.

Uniformity test based on discrepancy is shown to be more powerful than other alternatives [15]. Our strategy is to conduct uniformity testing based on simple statistics such as χ^2 first; if these simple tests fails to reject the null, then carry out the more complicated yet more powerful testing based on discrepancy. By noting that uniformity in $[a,b] = \prod_{j=1}^d [a_j, b_j]$ implies uniformity in each dimension; thus, we divide jth dimension into m bins $[a_j, a_j + (b_j - a_j)/m, ..., [a_j + (b_j - a_j)(m-2)/m, a_j + (b_j - a_j)(m-1)/m]$ and test the multinomial uniformity for the counts in bins, there are total d tests. We also keep track of the gaps at $a_j + (b_j - a_j)/m, ..., a_j + (b_j - a_j)(m-1)/m$, where the

- *Kun Yang is a PhD student of Stanford University at Wong Lab. E-mail: kunyang@stanford.edu.*
- *Wing Hung Wong is Stephen R. Pierce Family Goldman Sachs Professor in Science and Human Health, Professor of Statistics and Professor of Health Research and Policy at Stanford University. E-mail: whwong@stanford.edu.*

gap g_{jk} is defined as

$$g_{jk} = |\frac{1}{n}\sum_{i=1}^{n}\mathbf{1}(x_{ij} < a_j + (b_j - a_j)k/m) - \frac{k}{m}|$$

for $k = 1, ..., (m-1)$, there are total $(m-1)d$ gaps recorded. If any of the d tests reject the null, then we split the cube into two sub-cubes along the dimension and location corresponding to maximum gap (Figure 1); otherwise, we carry out the discrepancy based uniformity test and split the cube if it is rejected. For each cube, we carry out the tests and split it as above recursively until the null can not be rejected.

As detailed in Materials and Methods section of Appenidx, the output of the density estimation is a binary partition of the sample space with associated density in each sub-region. The density, which is a piecewise constant function, is

$$\hat{p}(x) = \sum_{i=1}^{n} d(r_i)\mathbf{1}(x \in r_i) \quad (1)$$

where $\mathbf{1}$ is indicator function; $\{r_i, d(r_i)\}_{i=1}^{n}$ is a list of pairs of sub-regions and corresponding densities (Figure 2). Note that the number of sub-regions is usually far less than the data size, hence $\hat{p}(x)$ is a concise summary of the data. For a given binary tree with the partition information encoded in each node, one can uniquely map it to a split of the sample space by recursive tree traversal. This one-to-one correspondence motivates us to utilize it as a proxy to visualize and manipulate the origin partition in high dimensions. We call this class of binary trees as "partition trees".

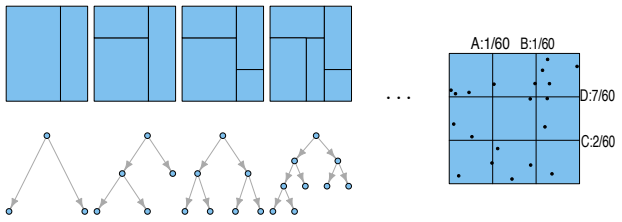

Fig. 1. Left: A sequence of binary partitions in two dimensional cube and the corresponding partition trees. From left to right, $t = 2, 3, 4, 5$. More information can be encoded in nodes, e.g., the dimensions and locations where the splits occur. Right: the gaps with $m = 3$, we split the cube at location D if the hypothesis is rejected.

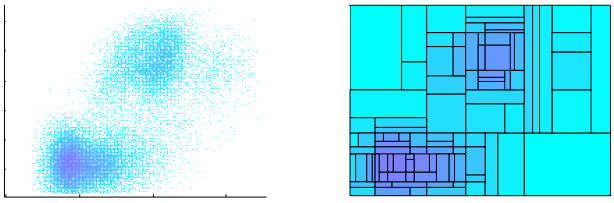

Fig. 2. An illustration of $p(x)$ in (1) in 2 dimension. Left: the point cloud; Right: the learned partition and associated densities displayed by a colormap.

3 Sub-level Tree

The sub-level tree (SLT) or tree of sub-level sets is widely used to represent energy landscapes [21]. An SLT summarizes the hierarchy among various local maxima and minima in the configuration space. The topology of an SLT is a tree and each inner node on the tree is a critical level that connects two or more separate regions in the domain. Given a density function $p(x)$ on Ω, let

$$\Omega_\eta^p = \{x : p(x) \geq \eta\}$$

be the sub-level set at level η and $conn(\Omega_\eta^p)$ be the set of connected components. One can verify the following properties.

Property. *For any* $0 \leq \eta' < \eta$,

1 $\forall X \in conn(\Omega_\eta^p)$, $\exists X' \in conn(\Omega_{\eta'}^p)$ *such that* $X \subseteq X'$ *and* X' *is defined as the parent of* X.

2 $\forall X \in conn(\Omega_\eta^p)$ *and* $X' \in conn(\Omega_{\eta'}^p)$, *either* $X \subseteq X'$ *or* $X \cap X' = \emptyset$

For a sequence $0 \leq \eta_1 < \eta_2 < \cdots < \eta_l$ and $\#conn(\Omega_{\eta_1}^p) = 1$, the above property de facto provides an algorithm to construct SLT. As illustrated in Figure 3: $\#conn(\Omega_\eta^p) = 1$ when $\eta < \eta_E$; Ω_η^p branches into two components when $\eta \in [\eta_E, \eta_D]$; Ω_η^p has one component again when $\eta \in (\eta_D, \eta_C)$; Ω_η^p splits into two smaller components at $\eta \in [\eta_C, \eta_B]$ and shrinks into one at $\eta \in (\eta_B, \eta_A]$. The corresponding SLT is constructed according to the parental relation defined in Property 1.

With the density estimate at hand, one may construct SLT for samples instead of a given energy or density function. Unlike kernel density estimation that suffers from many local bumps and results in an overly complicated SLT, piecewise constant function $\hat{p}(x)$ is well suited for this purpose, partially because it smoothes out the minor fluctuations and takes only limited number of values; moreover, its simple structure makes the construction of such graph easy; according to the algorithm in A.3, each sub-region of $\hat{p}(x)$ becomes a node on SLT. This tree representation has merits in several aspects: i) SLT provides a tree visualization of the data, which is especially useful when the data are multidimensional; ii) SLT provides a sparse summary of the scattering of the data, which is useful to detect modes ("mode seeking" is an important and widely used technique in computer vision [5]) and patterns in large data sets; iii) SLT is a high level abstraction of the data and can be use to extract new features.

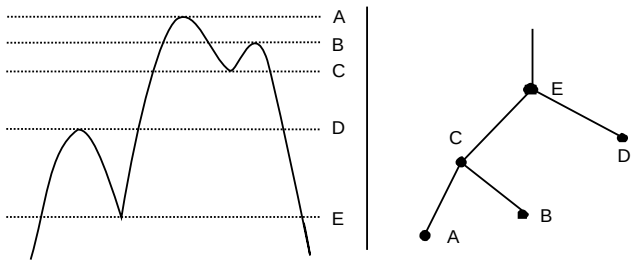

Fig. 3. A hypothetical density function (left) and its sub-level tree (right).

The algorithms to construct the binary partitioned piecewise density function and to build "partition tree" and "sub-level tree" are given in Materials and Methods.

4 Results

We first use a simulation to illustrate the basic method and demonstrate its properties such as the invariance to rotation and translation, then we apply them to two different kinds of data, namely flow cytometry data from mouse bone marrow and network data, in each case discovering the relevant hierarchies.

4.1 Simulations

Consider a Gaussian mixture:

$$p(x) = \sum_{i=1}^{4} \pi_i \mathcal{N}(\mu_i, \Sigma) \quad (2)$$

where $(\pi_1, \pi_2, \pi_3, \pi_4) = (.25, .25, .25, .25)$ and

$$\begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix} = \begin{pmatrix} 2 & 2 & & \cdots \\ -2 & 2 & & \cdots \\ & -2 & 2 & \cdots \\ & -2 & -2 & \cdots \end{pmatrix}_{4\times 10}$$

and

$$\Sigma = \begin{pmatrix} 1 & .1 & & & & \cdots \\ .1 & 1 & .1 & & & \cdots \\ & .1 & 1 & .1 & & \cdots \\ & & .1 & 1 & .1 & \cdots \\ & & & .1 & 1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ & & & .1 & 1 & .1 \\ & & & & .1 & 1 \end{pmatrix}_{10\times 10}$$

where void entries are 0s. In a generative model perspective [2], the data generation process can be represented schematically as in Figure 4: the cluster index is sampled according to π, then x is sampled from corresponding Gaussian distribution.

50,000 samples are sampled from (2), we use our methods to let the data "speak" for itself, i.e., to recover the hierarchy in Figure 4. The partition tree and sub-level tree are shown in Figure 5.ab. It is clear that the four branches of the sub-level tree in Figure 5.b correspond to the four clusters in $p(x)$. Moreover, richer information is contained in the trees: the two sub-branches correspond to the fact that cluster 1, 2 and cluster 3, 4 are closer to each other, so they merge before the four clusters becoming one; in fact, only the sub-branches are visible if the highest 5 levels of partition tree are trimmed down as shown in Figure 5.d. Figure 5.c demonstrates the invariant of sub-level tree under rotation and translation. In a word, without knowing the distribution a priori, the hierarchy in the data is revealed by our methods.

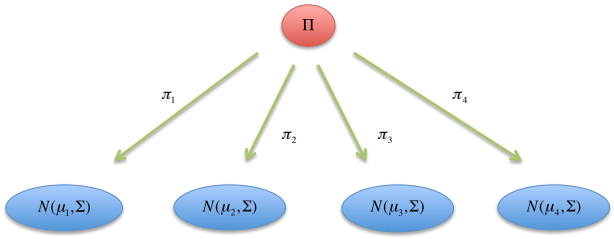

Fig. 4. A schematic representation of Gaussian Mixture from a generative model perspective.

4.2 Real Data

4.2.1 Flow Cytometry

Multi-parameter flow cytometry allows to measure multiple characteristics of single cells simultaneously; it provides insights into cellular differentiation, cellular hierarchy and disease diagnostics. Despite the increase in throughput and the number of parameters per single cell, there are limited number of methods for visualizing and analyzing multidimensional single-cell data. Moreover, cell differentiation creates the underlying hierarchy among the cell populations; traditional clustering algorithms are capable of finding mature cell populations (heterogeneity), whereas they ignore the continuity of phenotypes. As an attempt to capture this important aspect in cell populations, we apply our methods to the mouse bone marrow data studied in [17].

We regard each cell as one observation in the sample space, i.e., if there are n markers attached to a single cell, then the whole data set is generated from a hypothetical n dimensional distribution. Mature cell populations concentrate in some high density areas, i.e., the modes or local maxima on the domain. Through learning the n dimensional density and constructing the affiliated sub-level trees by adjusting the levels of partition tree, the cell populations are clustered around branches on the sub-level tree; based on the expression levels of markers in these populations, we can infer their hierarchy accordingly.

One practical issue needs to be addressed for most of the Cytometry analysis techniques: there is asymmetry in sub-populations; by

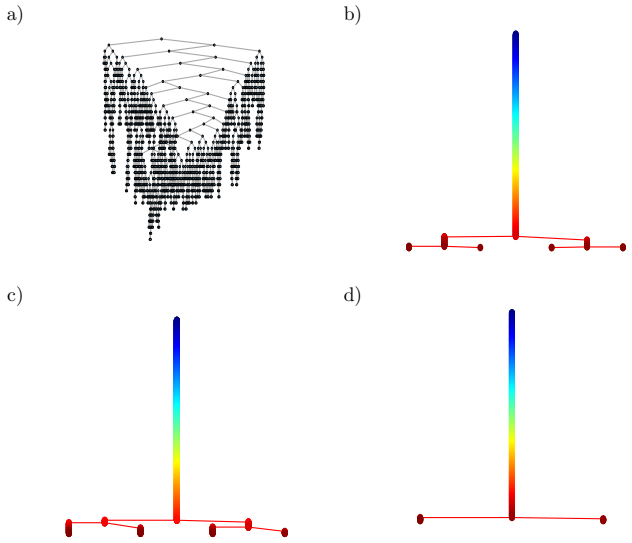

Fig. 5. Partition tree and sub-level trees for samples generated from the Guassian Mixture, the colors from blue to red on SLT represents the average densities from low to high as defined in (5): a) partition tree; b) corresponding sub-level tree; c) sub-level tree of the rotated and translated samples; d) sub-level tree after trimming down the highest 5 levels.

optimizing a predefined loss function, it is possible that some sparse yet crucial populations are overlooked if the algorithms take most of the efforts to control the loss in denser areas. A remedy for this issue is to perform a down-sampling [1, 17] step to roughly equalize the densities among populations then up-sampling after populations are identified; however, this step is dangerous that it may fails to sample enough cells in sparse populations, as a result, these populations are lost in the down-sampled data. In contrast, our approach does not require down-sampling step, and the asymmetry among populations are captured by their densities.

For the mouse bone marrow data, we choose the 8 markers (SSA-C, CD11b, B220, TCR-β, CD4, CD8, c-kit, Sca-1) that are relevant to the cell types of interests; the number of cells is ∼380,000 after removing mutli-cell aggregates and co-incident events. As shown in Figure 6, 13 sub-populations are identified ([17] and its supplementary materials); we arrange them into a hierarchical dendrogram: at first level, they are grouped by expression levels of CD11b; subsequently, the CD11b- sub-populations are grouped according to B220 and TCR-b then further splitted according to CD4 and CD8 on the next level; the CD11b+ sub-populations are grouped by B220 then by TCR-b.

4.2.2 Community structure in Social Networks

Diverse systems in various fields take the form of networks. In this study, we consider the community property which is found in many real networks such as social networks, bio networks and technological networks. In this example, we offer another approach to visualize the structure of the network by our sub-level tree algorithm. Analogues to the hierarchical clustering [20], it is a tree representation; however, it is much sparser and reveals the communities in its branches. We demonstrate that our methods can be used to detect the communities and reveal their denseness (cohesiveness) and discover the "transitional" nodes between the communities.

Given an undirected, unweighted $n-$vertices graph (network) $G = (V,E)$. The Laplacian matrix is defined as

$$L_{ij} = \begin{cases} 1, & i \sim j \\ -d_i, & i = j \\ 0, & i \nsim j \end{cases} \tag{3}$$

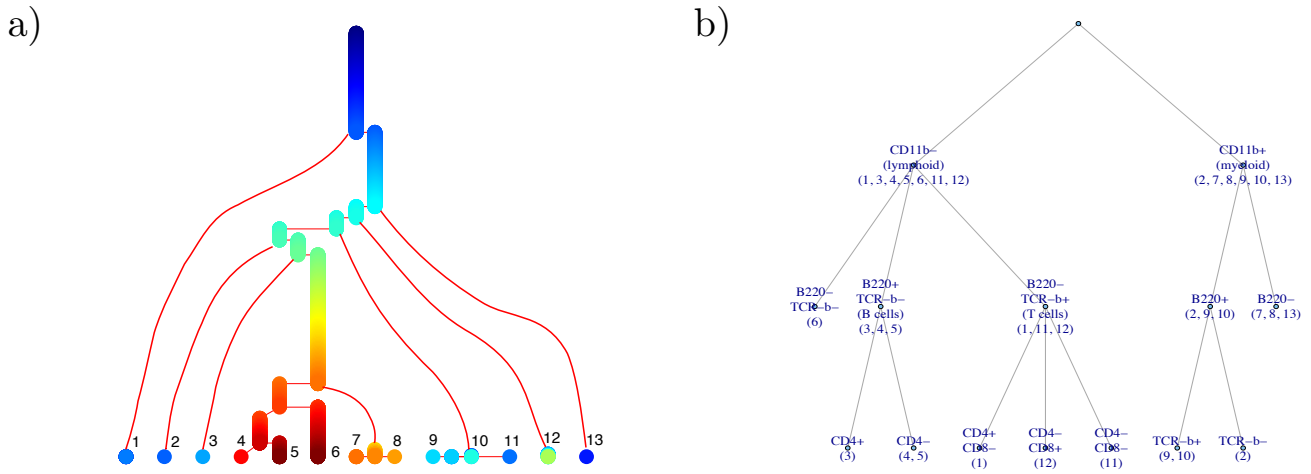

Fig. 6. Mouse bone marrow: a) sub-level tree learned from 8 markers, CD11b, B220, TCR-β, CD4, CD8, c-kit, Sca-1; b) corresponding cellular hierarchy built from the expression levels of markers in each sub-populations according to the marker sequence: CD11b, B220, TCR-β, CD4, CD8.

where $i \sim j$ ($i \nsim j$) means that the ith and jth vertices are (not) adjacent, and d_i is the degree of the vertex. In the spectral methods of graph clustering [16], we select the leading k eigenvectors (or excluding the first one[1]) of L (or regularized L [3]) and apply the $k-$means clustering algorithm; since clusters found by $k-$means are related to the modes of the underlying distribution and sub-level tree can also capture these modes, we replace $k-$means step by sub-level tree algorithm on these n $k-$dim vectors instead. The vertices belonging to sub-regions on a branch of SLT correspond to a community. However, some sub-regions have low densities and are added to S (S is defined in A.3) after the merging of branches according to our algorithm A.3. Each vertex in such sub-regions is in either of the two scenarios: 1) all its edges are connecting only one community, we assign it back to that community; 2) it serves as a hub that connects different communities, we define it as a "transitional" vertex since it plays a key role in the formation of communities.

We simulate a 1,000 vertices network and define the adjacency matrix M as follows: 1) Assign $M_{i,i+1} = 1, i = 1,2,...,999$ to make the network connected; 2) Construct three communities: $A = \{1,...,300\}$, $B = \{301,...,600\}$, $C = \{601,...,1,000\}$ with the edges in each community assigned as: i) $M_{i,j} = 1, i,j \in A$ with probability 0.01; ii) $M_{i,j} = 1, i,j \in B$ with probability 0.02; iii) $M_{i,j} = 1, i,j \in C$ with probability 0.008; and the edges between communities assigned as: i) $M_{i,j} = 1, i \in A, j \in B$ with probability 0.0001; ii) $M_{i,j} = 1, i \in B, j \in C$ with probability 0.0001; iii) $M_{i,j} = 1, i \in C, j \in A$ with probability 0.0005. We use the 3 leading eigenvectors of L to learn a binary partition and sub-level tree. In Figure 7, the three communities are identified on the branches of SLT; since A and C are closer to each other, their corresponding branches on SLT merge first.

We also apply our methods to classic dolphin social network [12, 11], it was constructed from observations of a community of 62 bottlenose dolphins over a period of seven years between 1994 and 2001. The two communities are correctly identified as shown in Figure 8.c and the relative "densities"(cohesiveness) of both communities are also colored in Figure 8.b. More interestingly, SN100, the individual with the highest connectivity in both communities and playing an important role in the fission and reunion of the dolphin community, are identified as a "transitional" vertex.

[1]It is obvious that $\overline{(1,1,...,1)}$ is the first eigenvector for any L and is trivial and uninteresting [13], we can omit it.

5 Discussion

Complex data can be understood in different perspectives. Classic methods display simple statistics such as mean, variance or point clouds with dimension no more than 3; early attempts to visualize high dimensional data, such as Chernoff faces [4], are applicable to relatively small data sets; as the data size increases, current approaches concentrate on the sparse representation which manages to display the information as much as possible, e.g., [10] tries to capture the topology ("shape") of the data and summarize the information in a graph.

Our methods are designed to mine another aspect of the data—hierarchy. They are non-parametric and unsupervised in nature, thus they do not suffer from the bias of specific model or assumptions. In Results, we show that they are applicable to different types of problems and capable of extracting the hierarchical information. The methods are generally useful to various problems after necessary processing steps on the raw data, for example, in dealing with real data, it is possible that the data defined in the original parameter space are not immediately useful and some data transformation is important.

A Materials and Methods

A.1 Binary Partition by Discrepancy

$P(\cdot)$ define the points and $\Pr(\cdot)$ define the probability mass in a sub-region respectively. Without loss of generality, we assume that $\Omega = [0,1)^d$ and $P(\Omega) = \{x_i = (x_{i1}, x_{i2}, ..., x_{id})\}_{i=1}^n$.

procedure DENSITY-ESTIMATOR(Ω, m, α)
  $T = \{\Omega\}$, $\Pr(\Omega) = 1$
  while true **do**
    $\tilde{T} = \emptyset$
    for each $r = \prod_{j=1}^d [a_j, b_j] \in T$ **do**
      scale $P(r) = \{x_{r_j}\}_{j=1}^{n_r}$
      to $\tilde{P}(r) = \{\tilde{x}_{r_j} = (\frac{x_{r_j,1} - a_1}{b_1}, ..., \frac{x_{r_j,d} - a_d}{b_d})\}_{j=1}^{n_r}$
      for $j \leftarrow 1, ..., d$ **do**
        $B_i = 0, i = 1,2,...,m$
        for each $x = (x_1, ..., x_d) \in \tilde{P}(r)$ **do**
          $B_{\lfloor m x_j \rfloor + 1} = B_{\lfloor m x_j \rfloor + 1} + 1$
          ▷ count the number of points in each bin
        test the multinomial uniformity of $\{B_1, ..., B_m\}$
        record the gaps as $\{g_{jk}\}_{k=1}^{m-1}$
      if any of the d tests are rejected **then**
        split r into $\{r_1, r_2\}$ along the maximum gap $\{g_{jk}\}$
      else

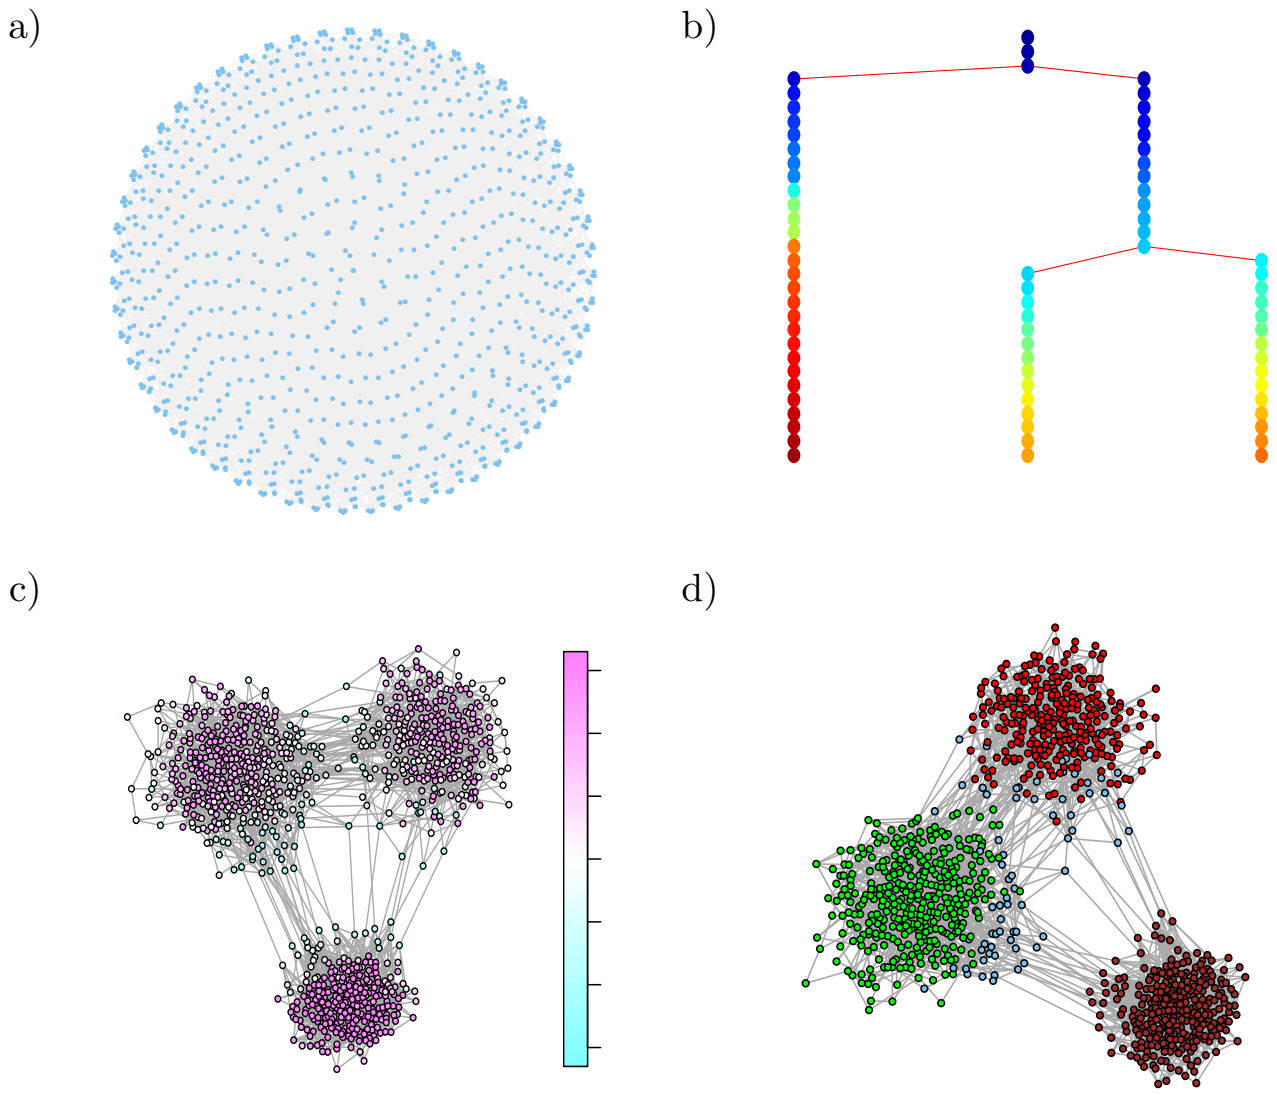

Fig. 7. Network: a) original network plotted on the sphere; b) corresponding sub-level tree of the 3 leading eigenvectors of the Laplacian; the two closer communities merge first, then merge with the third one; c) vertices in the network are colored according to densities; d) communities colored by red, brown and green, the transitional vertices are colored by blue.

18: test the uniformity by discrepancy
19: split r as in line 16 if it is rejected
20: **if** r is divided **then**
21: $\tilde{T} = \tilde{T} \cup \{r_1, r_2\}$
22: $\Pr(r_1) = \Pr(r)\frac{\#P(r_1)+\alpha}{\#P(r)+2\alpha}$
23: $\Pr(r_2) = \Pr(r) - \Pr(r_1)$
24: **else**
25: $\tilde{T} = \tilde{T} \cup \{r\}$
26: **if** $\tilde{T} == T$ **then**
27: **return** T
28: **else**
29: $T = \tilde{T}$

Remark A.1. The partition tree can be constructed as a byproduct by bookkeeping the parental relations of partitions.

The density in r is recovered by $d(r) = \Pr(r)/|r|$, where $|r|$ is the volume of r; $\alpha > 0$ is a Laplace smoother (pseudo count). In line 13, Pearson's Chi-square test or log likelihood test can be used; in line 18, we apply the symmetric discrepancy based test [9] as follows: let

$$A = \frac{1}{n}\sum_{i=1}^{n}\prod_{j=1}^{d}(1+2x_{ij}-2x_{ij}^2)$$

$$B = \frac{2^{d-1}}{n(n-1)}\sum_{i<j}\prod_{k=1}^{d}(1-|x_{ik}-x_{jk}|)$$

$$C = (4/3)^d, \eta = (9/5)^d - (6/9)^d$$

then

$$\sqrt{n}[(A-C)+2(B-C)]/(5\sqrt{\eta}) \xrightarrow{\mathscr{D}} \mathscr{N}(0,1) \quad (4)$$

Note that B can be computed in $O(n\log^{d-1} n)$ according to Frank and Heinrich's algorithm [6]; at level t, the number of samples in each sub-region is $n_i, i = 1,...,t$, the complexity is

$$\sum_{i=1}^{t} n_i \log^{d-1} n_i \le \sum_{i=1}^{t} n_i \log^{d-1} n = n\log^{d-1} n$$

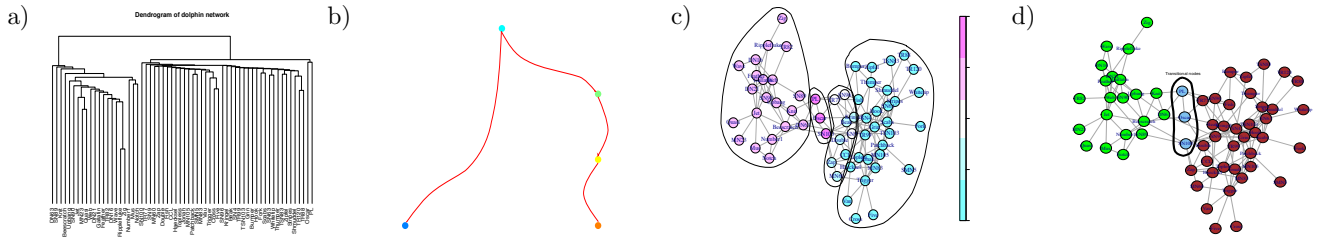

Fig. 8. Dolphin Social Network: a) the dendrogram by hierarchical clustering according to [20]; b) sub-level tree of the leading 2 eigenvectors of the Laplacian; c) the dolphin network that vertices are colored and grouped according to densities; d) two dolphin communities and the "transitional" vertices, particularly, SN100, which plays an important role in the formation of communities, is identified.

thus, the total complexity is $O(l \cdot n \log^{d-1} n)$, where l is the deepest level, which is a moderate number in our experience and can be specified by the user as well.

As shown in [9], discrepancy based test is powerful even when the sample size is less than 1,000. We can also compute (4) by subsampling a smaller sample (say 500 points, which works very well in our examples); since there are t sub-regions in level t and the uniformity test in each sub-region takes $O(m \log^{d-1} m)$ with m sub-samples, so the complexity is at most

$$\sum_{t=1}^{l} O(m \log^{d-1} m) t = O(m \log^{d-1} m \cdot l^2)$$

Both computing strategies yield similar empirical results, but the $O(m \log^{d-1} m \cdot l^2)$ one becomes attractive when the data size is large.

A.2 Graph of the Partition

For a given partition P and the list of pairs of sub-regions and corresponding densities $\{r_i, d(r_i)\}_{i=1}^n$ as in (1), we build a graph G based on the adjacency of sub-regions and each sub-region is a node on the graph. The algorithm to determine the adjacency of sub-region i, j is:

procedure IS-ADJACENT(r_i, r_j)
  $c_k = (c_{k1}, ..., c_{kd})$: the center of $r_k, k \in \{i, j\}$
  $l_k = (l_{k1}, ..., l_{kd})$: the width of r_k in each dimension, $k \in \{i, j\}$
  for $k \leftarrow 1, ..., d$ **do**
    if $|c_{ik} - c_{jk}| > (l_{ik} + l_{jk})/2$ **then**
      return False
  return True

G is constructed by connecting adjacent sub-regions. When G has $k > 1$ connected components $\{c_1, c_2, ..., c_k\}$, a "virtual region" r is added into G with density zero and it connects the region in each $\{c_i\}_{i=1}^k$ with lowest density in order to make G connected.

A.3 Sub-level Tree

Starting with empty set S_0 at step 0, the sub-region is added into S sequentially according to the decreasing order of densities. At ith step, we compute the connected components based on the induced subgraph $G(S_i)$. Suppose $S_i = \{g_1, g_2, ..., g_j\}$, where g_i is the connected components and j is the number of components; at step $i+1$, there are two scenarios when r is added into S_i: i) r is adjacent to multiple components $g_{k_1}, ..., g_{k_l}$, then $S_{i+1} = \{S_i \backslash \{g_{k_i}\}_{i=1}^l, r \cup \{g_{k_i}\}_{i=1}^l\}$ and r is the parent of latest added sub-region in each g_{k_i}; ii) r is disconnected with all the components, then $S_{i+1} = \{S_i, r\}$ and r is a leaf. The complete description of the algorithm is:

procedure SUB-LEVEL-TREE(P)
  Build G from P
  $S_0 = \emptyset$
  $R_0 = \emptyset$
  for $k \leftarrow 1, n$ **do** ▷ n is the number of nodes in G
    $S_{k-1} = \{g_1, g_2, ..., g_l\}$
      ▷ $\{g_i\}_{i=1}^l$ are the connected components in $G(S_{k-1})$
    $R_{k-1} = \{r_1, r_2, ..., r_l\}$
      ▷ r_i is the last sub-region added into $g_i, i = 1, ..., l$
    if $r_{(k)}$ is adjacent to $\{g_{i_1}, g_{i_2}, ..., g_{i_m}\}$ **then**
      ▷ $r_{(k)}$ has kth largest density in P
      $S_k = \{r_{(k)} \cup \{g_{i_j}\}_{j=1}^m, S_{k-1} \backslash \cup \{g_{i_j}\}_{j=1}^m\}$
      $R_k = \{r_{(k)}, R_{k-1} \backslash \cup \{r_{i_j}\}_{j=1}^m\}$
      $\wp(r_{i_j}) = r_{(k)}, j = 1, ..., m$
        ▷ $\wp$ stores the parent of each sub-region
      Color($r_{(k)}$) = average density($r_{(k)} \cup \{g_{i_j}\}_{j=1}^m$)
    else
      $S_k = \{S_{k-1}, r_{(k)}\}$
      $R_k = \{R_{k-1}, r_{(k)}\}$
      Color($r_{(k)}$) = average density($r_{(k)}$)
  return $\wp$, Color

At each step, we also keep track of the average density in each component; the average density is defined as the ratio between the total mass and total volume in the component, i.e., the average density of g is

$$\text{average density}(g) = \frac{\sum_{r \in g} |r| d(r)}{\sum_{r \in g} |r|} \tag{5}$$

The tree nodes can be colored according to the average density when the sub-region is included in S for the first time.

A.4 Software

A Matlab implementation is available upon request.

Acknowledgments

Kun Yang is supported by General Wang Yaowu Stanford Graduate Fellowship and The Simons Math+X fellowship; Wing Hung Wong is supported by NSF grants DMS 0906044 and 1330132.

References

[1] N. Aghaeepour, G. Finak, H. Hoos, T. R. Mosmann, R. Brinkman, R. Gottardo, R. H. Scheuermann, F. Consortium, and D. Consortium. Critical assessment of automated flow cytometry data analysis techniques. *Nature methods*, 2013.

[2] C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.

[3] K. Chaudhuri, F. C. Graham, and A. Tsiatas. Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research-Proceedings Track*, 23:35.1–35.23, 2012.

[4] H. Chernoff. The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68(342):361–368, 1973.

[5] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(5):603–619, 2002.

[6] C. Doerr, M. Gnewuch, and M. Wahlström. Calculation of discrepancy measures and applications. *Preprint*, 2013.

[7] L. C. Freeman. Visualizing social networks. *Journal of social structure*, 1(1):4, 2000.

[8] F. Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of Computation of the American Mathematical Society*, 67(221):299–322, 1998.

[9] J.-J. Liang, K.-T. Fang, F. Hickernell, and R. Li. Testing multivariate uniformity and its applications. *Mathematics of Computation*, 70(233):337–355, 2001.

[10] P. Lum, G. Singh, A. Lehman, T. Ishkanov, M. Vejdemo-Johansson, M. Alagappan, J. Carlsson, and G. Carlsson. Extracting insights from the shape of complex data using topology. *Scientific reports*, 3, 2013.

[11] D. Lusseau and M. E. Newman. Identifying the role that animals play in their social networks. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(Suppl 6):S477–S481, 2004.

[12] D. Lusseau, K. Schneider, O. J. Boisseau, P. Haase, E. Slooten, and S. M. Dawson. The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations. *Behavioral Ecology and Sociobiology*, 54(4):396–405, 2003.

[13] M. E. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104, 2006.

[14] A. B. Owen. Quasi-monte carlo sampling. *Monte Carlo Ray Tracing: Siggraph*, pages 69–88, 2003.

[15] A. Petrie and T. R. Willemain. An empirical study of tests for uniformity in multidimensional data. *Computational Statistics & Data Analysis*, 2013.

[16] T. Qin and K. Rohe. Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *Advances in Neural Information Processing Systems*, pages 3120–3128, 2013.

[17] P. Qiu, E. F. Simonds, S. C. Bendall, K. D. Gibbs Jr, R. V. Bruggner, M. D. Linderman, K. Sachs, G. P. Nolan, and S. K. Plevritis. Extracting a cellular hierarchy from high-dimensional cytometry data with spade. *Nature biotechnology*, 29(10):886–891, 2011.

[18] H. van den Bedem, G. Bhabha, K. Yang, P. E. Wright, and J. S. Fraser. Automated identification of functional dynamic contact networks from x-ray crystallography. *Nature methods*, 10(9):896–902, 2013.

[19] M. A. Wilson. Visualizing networks of mobility in proteins. *Nature methods*, 10(9):835–837, 2013.

[20] S. Zhang, X.-M. Ning, and X.-S. Zhang. Graph kernels, hierarchical clustering, and network community structure: experiments and comparative analysis. *The European Physical Journal B*, 57(1):67–74, 2007.

[21] Q. Zhou and W. H. Wong. Energy landscape of a spin-glass model: Exploration and characterization. *Physical Review E*, 79(5):051117, 2009.